

Programming the I²C interface

Author: Mitchell Kahn

Dr. Dobb's

JOURNAL

The Inter-Integrated Circuit Bus ("I²C Bus" for short) is a two-wire, synchronous, serial interface designed primarily for communication between intelligent IC devices. The I²C bus offers several advantages over "traditional" serial interfaces such as Microwire and RS-232. Among the advanced features of I²C are multimaster operation, automatic baud-rate adjustment, and "plug-and-play" network extensions.

Mention the I²C bus to a group of American engineers and you'll likely get hit with an abundance of blank stares. I say American engineers because until recently the I²C bus was primarily a European phenomenon. Within the last year, however, interest in I²C in the United States has risen dramatically. Embedded systems designers are realizing the cost, space, and power savings afforded by robust serial interchip protocols.

The idea of serial interconnect between integrated circuits is not new. Many semiconductor vendors offer devices designed to "talk" via serial links with other processors. Current examples include Microwire (National Semiconductor), SPI (Motorola), and most recently Echelon's Neuron chips. In all cases, the goal is the same: to reduce the wiring and pincount necessary for a parallel data bus. It simply does not make

Mitch is a senior strategic development engineer for Intel and can be contacted at 5000 W. Chandler Blvd., Chandler, AZ 85226 or at mkahn@sedona.intel.com.



economic sense to route a full-speed parallel bus to a slow peripheral.

Unfortunately for most serial-bus-capable devices, the choice of a bus protocol will dictate the CPU architecture. For example, only two CPU architectures implement an on-chip I²C port. If your choice of architecture precludes use of these architectures, then your only option is to implement the protocol in software.

The software implementation of the I²C protocol discussed in this article came about as a result of an implicit challenge during a staff meeting. One of our managers proposed that we hire a consultant to write a software I²C driver for the Intel 80C186EB embedded processor. Being somewhat new to the

group, I took exception (although not verbally!) to his suggestion. A weekend of intense hacking later, I presented the first prototype of the driver. My reward? I got to write a generic version of the driver for general distribution.

Design Trade-offs

Three distinct tasks are involved in implementing the I²C protocol: watching the bus, waiting for a specific amount of time, and driving the bus. This became apparent when I flowcharted a byte of a typical bus transaction; see Figure 1. The time delays associated with creating the bus waveforms would normally have been relegated to the 80C186EB's on-chip timers. I could not, however, assume that the end users of my code would be able to spare a timer for the software I²C port. I had to forego the elegance (and to some extent accuracy) of the on-chip timers for the sledgehammer approach of software timing loops. Luckily, the I²C protocol is extremely forgiving with regard to timing accuracy. The decision to use assembly instead of a high-level language stemmed directly from the need to control program-execution time. I had neither the time nor the inclination to hand-tune high-level code.

Having made the decision to use assembly language, I faced my next problem: Could I make the code portable? Intel offers a plethora of CPU and embedded-controller architectures. Would it be possible to make the code somewhat portable between disparate assembly languages? I found my answer in the use of macros.

Programming the I²C interface

All the basic building blocks of the I²C protocol (watching, waiting, and doing) can be compartmentalized into distinct macros. The algorithms that make up the I²C driver are written with these macros as the framework. You don't need to understand the intricacies of the I²C protocol to port these routines—you just need to know how to make your CPU watch, wait, and do.

For example, a 4.7 μ s delay is a common event during a transfer. The macro %Wait_4_7_us implements just such a delay by using the 8086 LOOP instruction with a couple of NOPs for tuning; see Example 1(a). Total execution time is readily calculated from instruction timing tables. The same macro is ported to the i960 architecture in Example 1(b). Although I am a neophyte when it

comes to i960 programming, I had no problems porting the core macros.

Hardware Dependencies

A few words about the target hardware are in order before I discuss the code. Any implementation of the I²C protocol requires two open-drain (or open-collector), bidirectional port pins for the Serial Clock (SCL) and Serial Data (SDA) lines. The code in this article was designed for the 80C186EB embedded processor, which has two open-drain ports on-chip. The two pins, P2.6 (SCL) and P2.7 (SDA), are part of a larger 8-bit port. Processors without open-drain I/O ports can easily implement I²C with the addition of an external open-collector latch.

Two special-function registers, P2PIN and P2LTCH, are used to read and write the state of the port pins. The 80C186EB allows the special-function registers to be located anywhere in either memory or I/O space. For this implementation, I chose to leave the registers in I/O space, even though this limited my choice of instructions. The 80186 architecture does not provide for read-modify-write instructions in I/O space (an AND to I/O, for example); it can only load and store (IN and OUT). So why did I limit myself? Again, I had to assume the lowest common denominator for our customers when designing my code.

Building the Framework

Early on in development, I decided to partition my code macros according to physical processes involved in the I²C

protocol. Code not directly involved in mimicking the actions of a hardware I²C port was not written as macros. For example, the code necessary to access the stack frame is not written as a macro, whereas the code needed to toggle the clock line is. This was done to isolate architecture-dependent code sequences from the more generic I²C functions. Macros were also not used for "gray areas" such as the shifting of serial data, which is both architecture dependent and physical in nature. The I²C functions that passed the litmus test fell into the three aforementioned categories of watching, waiting, and doing.

The "waiting" macros provide a fixed-minimum time delay. They are implemented using a simple LOOP \$ delay. The LOOP instruction decrements the CX register, then branches to the target (in this case itself) if the result is non-zero. The delay is $(n-1) \times 15 + 5$ clocks, where n is the starting value in the CX register. All the delays were calculated assuming a 16-MHz clock rate (62.5 nanoseconds per clock). The code still works at lower CPU speeds because the I²C protocol only specifies minimum timings. In fact, the delay macros are only "accurate enough," providing timings as close as I could get to the specified minimum without undue tuning.

The "watching" macros are "spin-on-bit" polling loops. These pieces of code wait for a transition on the appropriate I²C line to occur before allowing execution to continue. There are two polling macros for each of the two I²C signal lines: one for high-to-low transitions and one for low-to-high transitions. The




Figure 1: Flowchart of process for transmission of a single bit.

```

(a)
%DEFINE(Wait_4_7_us)
    mov     cx, 5          ; 4 clocks
    loop    $              ; 4*15+5 = 65 clocks
    nop     $              ; 3 clocks
    nop     $              ; 3 clocks
    nop     $              ; 3 clocks
    ; total = 75 clocks
    ; 75 * 62.5ns = 4.69μs (close enough)

(b)
define(Wait_4_7_us,
    lda     0x17, r4        ; instruction may be issued in parallel
    ; so assume no clocks
    ; compare and decrement counter in r4
    ; if !=0 branch back (predic taken
    ; branch)
    ;
    ; The cmpdeco and bne.t together take 3
    ; clocks in parallel minimum.
    ;
    ; 0x17 (25 decimal) * 3 = 75 clocks
    ; at 16MHz this is 4.69μs
)
  
```

Example 1: (a) 80C186 implementation of 4.7 μ s wait macro; (b) 80960CA implementation of 4.7 μ s wait macro.

Programming the I²C interface

polling of the SCL line that gives rise to an important feature of I²C: automatic, bit-by-bit baud-rate adjustment. Any device on the I²C bus may hold the clock line low in order to stall the bus for more time (a serial wait state). The other devices on the bus are then forced to poll the SCL line until the slow device releases control of the clock.

The %Get_SDA_Bit macro also falls under the category of "watching." Its function is simply to return the state of the SDA line without waiting for a transition. %Get_SDA_Bit is used primarily to pull the serial data off the bus when the clock is valid.

The "doing" macros control the state of the clock and data lines. As with the polling macros, there are four types—one for each transition of the SCL or SDA lines. The "doing" macros are named to reflect the physical operations they perform. For example, %Drive_SCL_Low always drives the SCL line to a low state. %Release_SCL_High, on the other hand, relinquishes control of the SCL line, which may then be pulled high or driven low by another device on the bus. A read-modify-write operation is used for the bit manipulation so that the other 6 bits of Port 2 are not affected by the I²C operations.

Getting on the Bus

Three procedures were created using the macro framework. I'll describe only the master transmit (Listing One, page




Figure 2: Flowchart for I²C transmit procedure.

3

106) and master receive functions (Listing Two, page 108), as they represent the needs of most I²C users. The slave procedure is long and intricate and will not be described here.

An I²C master transmission proceeds as follows:

1. The master polls the bus to see if it is in use.
2. The master generates a start condition on the bus.
3. The master broadcasts the slave address and expects an acknowledge (ACK) from the addressed slave.
4. The master transmits 0 or more bytes of data, expecting an ACK following each byte.
5. The master generates a stop condition and releases the bus.

The stack frame for the master transmit procedure, I2CXA.A86, includes a far pointer to the message for transmission, the byte count for the message, and the slave address. Far pointers and far procedure calls are used in all the procedures. No attempt was made to conform to a specific high-level language calling convention, although such a conversion would be trivial. The procedures save only the state of the modified segment registers.

The master transmit procedure performs error checking on the passed parameters before attempting to send the message. The maximum message length is set at 64 Kbytes by the segmentation of the 80186 memory space. This restriction could be removed by including code to handle segment boundaries. The transmit procedure also checks the direction bit in the slave address to ensure that a reception was not erroneously indicated. Errors are reported back to the calling procedure through the AX register. (The exact code is in Listing One.)

The first step in sending a message is getting on the I²C bus. The macro %Check_For_Bus_Free simply polls the bus to determine if any transactions are in progress. If so, the transmit procedure aborts with the appropriate error code. If the bus is free, a start condition is generated. The start condition is defined as a high-to-low transition of SDA with SCL high followed by a 4.7_μs pause. These waveforms are easily generated with the %Drive_SDA_Low and %Wait_4_7_μs macros.

All communication on the I²C bus between the stop and start conditions, including addressing and data, takes place as an 8-bit data value followed by an acknowledge bit. This leads to the natural nested loop structure for the body of the procedure; see Figure 2.

The inner loop is responsible for transmitting the 8 bits of each data byte. Each transmitted bit generates the appropriate data (SDA) and clock (SCL) waveforms while checking for both serial wait states and potential bus collisions. A bus collision occurs when two masters attempt to gain control of the

*Three distinct tasks
are involved in
implementing the
I²C protocol:
watching the bus,
waiting for a specific
amount of time, and
driving the bus*

bus simultaneously. The I²C protocol handles collisions with the simple rule: "He who transmits the first 0 on the SDA line wins the bus." To ensure that we (the master transmit procedure) own the bus, the SDA line is checked whenever transmitting a 1. If a 0 is present, then a collision has occurred (because another master is pulling the line low), and the transfer must be aborted.

Control is turned over to the outer loop after the 8 bits of data (or address) have been transmitted. The outer loop immediately checks for an acknowledge from the addressed slave. The transfer is aborted if an acknowledge is not received. At the end of the ACK bit the message length counter is decremented. Control is returned to the inner loop if more data remains, otherwise a stop condition is generated and the master transmit procedure terminates.

Registers are used for intermediate result storage throughout the body of the procedure. For example, the AH register is used to hold the current value (either address or data) being shifted onto the SDA line. This eliminates the need for local data storage within the procedure.

On the Receiving End

The steps involved in an I²C master receive transaction are almost identical to those in transmission:

1. The master polls the bus to see if it is in use.
2. The master generates a start condi-

Dr. Dobb's Journal, June 1992

Programming the I²C interface

- tion on the bus.
3. The master broadcasts the slave address and expects an ACK from the addressed slave.
4. The master receives 0 or more bytes of data and sends an ACK to the slave after each byte. The master signals the last byte by not sending an ACK.
5. The master generates a stop condition and releases the bus.

A far pointer to the receive buffer is passed on the stack to the master receive procedure. The remainder of the parameters—slave address and message count—are identical between the two procedures. The received message length is fixed at 64 Kbytes, again because of segmentation. The error-checking, bus-availability sensing, and start-condition generation sections of the receive procedure are lifted verbatim from the transmit code.

The structure of the receive procedure differs slightly once the start con-

dition has been generated; see Figure 3. The slave address is transmitted using one iteration of the transmit procedure's outer loop. Control is passed to the receive loop once the slave acknowledges its address.

The receive loop structure is patterned after that of the transmit procedure. The inner loop controls the clocking of the SCL line and the shifting of the serial data off the SDA line into the CPU. Eight iterations of the inner loop are performed to receive each byte. The outer loop stores the received byte in the buffer, decrements the byte count, then sends an ACK to the slave. The last data byte is signalled by not sending an ACK.

Using the Procedures

Listing Three (page 110) shows a short program that uses both the master transmit and master receive procedures. The call to procedure I2C_XMIT displays the word "BUS-" on a four-character, seven-segment display controlled by the SAA1064 I²C compatible display driver. The time of day is read from the PCF8583 real-time clock by the call to procedure I2C_RECV.

Please note that interrupts must be disabled during the execution of both procedures. An interruption at an inopportune time (when the master is not in control of the clock) could cause the bus to hang. If you need to service interrupts periodically, then enable them only when the clock is driven low.

These procedures have been tested on a wide array of I²C devices ranging from serial EEPROMs to voice synthesizers. No compatibility problems have been seen to date.

Enhancing the Code

I've kicked around many ideas for enhancing the I²C procedures. You could,

*All the basic
building blocks of
the I²C protocol
(watching, waiting,
and doing) can be
compartmentalized
into distinct macros*




Figure 3: Flowchart for I²C receive procedure.

for example, replace the timing loops with timed interrupts. That way, the CPU could perform useful work during the pauses. Along the same lines, the pauses could be scheduled using a real-time kernel, again improving CPU throughput. Finally, you could add a high-level language calling structure.

The use of timed interrupts adds an order of magnitude to the complexity of the code, but would be worth it for high-performance, real-time systems.

Conclusion

I²C is not the only game in town when it comes to serial protocols. Hopefully, some of the techniques presented here will carry over into the development of other "simulated" serial protocols, such as those targeted at the home-automation market. Who knows, maybe someday a snippet of my code may find its way into a truly intelligent dishwasher. I'll be waiting....

References

I²C Bus Specification, Philips Corporation (undated).

DDJ

Reprinted with permission of Dr.

Dobb's Journal, 1992

Entire contents copyright © 1992
by M&T Publishing, Inc.
Unless otherwise noted on specific articles.
All rights reserved.

ABP
American Business Press

The
Audit
Bureau